Fast Direct Methods for Gaussian Processes and the Analysis of NASA *Kepler* Mission Data

Sivaram Ambikasaran, Daniel Foreman-Mackey, Leslie Greengard,
David W. Hogg, and Michael O'Neil

Abstract—A number of problems in probability and statistics can be addressed using the multivariate normal (or multivariate Gaussian) distribution. In the one-dimensional case, computing the probability for a given mean and variance simply requires the evaluation of the corresponding Gaussian density. In the n-dimensional setting, however, it requires the inversion of an $n \times n$ covariance matrix, C, as well as the evaluation of its determinant, $\det(C)$. In many cases, the covariance matrix is of the form $C = \sigma^2 I + K$, where K is computed using a specified kernel, which depends on the data and additional parameters (called hyperparameters in Gaussian process computations). The matrix C is typically dense, causing standard direct methods for inversion and determinant evaluation to require $\mathcal{O}(n^3)$ work. This cost is prohibitive for large-scale modeling. Here, we show that for the most commonly used covariance functions, the matrix C can be hierarchically factored into a product of block low-rank updates of the identity matrix, yielding an $\mathcal{O}(n \log^2 n)$ algorithm for inversion, as discussed in Ambikasaran and Darve, 2013. More importantly, we show that this factorization enables the evaluation of the determinant $\det(C)$, permitting the direct calculation of probabilities in high dimensions under fairly broad assumption about the kernel defining K. Our fast algorithm brings many problems in marginalization and the adaptation of hyperparameters within practical reach using a single CPU core. The combination of nearly optimal scaling in terms of problem size with high-performance computing resources will permit the modeling of previously intractable problems. We illustrate the performance of the scheme on standard covariance kernels, and apply it to a real data set obtained from the *Kepler* Mission.

Index Terms—Covariance function, Gaussian process, hierarchical off-diagonal low-rank, direct solver, fast multipole method, Bayesian analysis, likelihood, evidence

arXiv:1403.6015v1 [math.NA] 24 Mar 2014

1 Introduction

A common task in probability and statistics is the computation of the numerical value of the posterior probability of some parameters $\boldsymbol{\theta}$ (conditional on some data $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$) using a multivariate Gaussian distribution. This requires the evaluation of

$$p(\boldsymbol{\theta}|\boldsymbol{x}, \boldsymbol{y}) \propto \frac{1}{(2\pi)^{n/2} |\det(C(\boldsymbol{x}))|^{1/2}} e^{-\frac{1}{2}\boldsymbol{y}^t C^{-1}(\boldsymbol{x})\boldsymbol{y}}, \quad (1)$$

where $C(\boldsymbol{x})$ is an $n \times n$ symmetric, positive-definite *covariance* matrix. In the one-dimensional case, $C(\boldsymbol{x})$ is simply the scalar variance. Thus, computing the probability requires only the evaluation of the corresponding Gaussian. In the n-dimensional setting, however, $C(\boldsymbol{x})$ is typically dense, so that its inversion requires $O(n^3)$ work as does the evaluation of its determinant $\det(C)$. This cost is prohibitive for large n.

In many cases, the covariance matrix C is assumed to be of the form $C(\boldsymbol{x}) = \sigma^2 I + K(\boldsymbol{x})$, where $K_{ij}(\boldsymbol{x}) = k(x_i, x_j)$. The function $k(x_i, x_j)$ is called the covariance function or covariance kernel, which, in turn, can depend on additional parameters, $\boldsymbol{\theta}$. Covariance matrices of this form universally appear in regression and classification problems when using Gaussian process priors [43]. Because many covariance kernels are similar to those that arise in computational physics, a substantial body of work over the past decades has produced a host of relevant fast algorithms, first for the rapid application of matrices such as K [17], [21], [26], [27], [55], and more recently on their inversion [4], [7], [10], [14], [17], [25], [33], [38]. We do not seek to further review the literature here, except to note that it is still a very active area of research.

Using the approach outlined in [4], we will show below that, under suitable conditions, the matrix C can be hierarchically factored into a product of block low-rank updates of the identity matrix, yielding an $\mathcal{O}(n \log^2 n)$ algorithm for inversion. More importantly (and perhaps somewhat surprising), we show that our factorization enables the evaluation of the determinant, $\det(C)$, in $\mathcal{O}(n \log n)$ operations. Together, these permit the efficient direct calculation of probabilities in high dimensions. Previously existing methods for inversion and determinant evaluation were based on either rough

- *S. Ambikasaran is with the Courant Institute at New York University, New York, NY, e-mail: sivaram@cims.nyu.edu.*
- *D. Foreman-Mackey is with the Department of Physics at New York University, New York, NY, e-mail: danfm@nyu.edu.*
- *D. W. Hogg is with the Department of Physics at New York University, New York, NY, e-mail: david.hogg@nyu.edu.*
- *L. Greengard is with the Courant Institute at New York University, New York, NY and the Simons Center for Data Analysis at the Simons Foundation, New York, NY, e-mail: greengard@cims.nyu.edu.*
- *M. O'Neil is with the Courant Institute at New York University, New York, NY, e-mail: oneil@cims.nyu.edu.*

Research by S. Ambikasaran and M. O'Neil is supported in part by the Air Force Office of Scientific Research under NSSEFF Program Award FA9550-10-1-0180. Research by L. Greengard is supported in part by the Simons Foundation and the Air Force Office of Scientific Research under NSSEFF Program Award FA9550-10-1-0180. Research by D. Foreman-Mackey and D. W. Hogg is supported in part by NASA grant NNX12AI50G, NSF grant IIS-1124794, the Gordon and Betty Moore Foundation, and the Alfred P. Sloan Foundation.

approximation methods or iterative methods [8], [12], [15], [47], [48]. These schemes are particularly ill-suited for computing determinants. Although bounds exist for sufficiently random and diagonally dominant matrices, they are often inadequate in the general case [11].

We will focus, in this paper, on describing and applying our new methods for handling large-scale covariance matrices (dense and full- or high-rank) to the problem of regression under Gaussian process priors. Other applications, such as marginalization and adaptation of hyperparameters are relatively straightforward, and the computational bottlenecks of each are discussed briefly.

Gaussian processes are the tool of choice for many statistical inference or decision theory problems in machine learning and the physical sciences. They are ideal when requirements include flexibility for the modeling of continuous functions. However, applications are limited by the computational cost of determinant calculation. The determinant calculation is required for Gaussian process likelihood evaluations (i.e., computation of any actual value of the probability of the data under the covariance hyperparameters, or evidence). This cost is even greater when the likelihood evaluation is placed *inside* an outer optimization or Markov chain Monte Carlo (MCMC) sampling loop. Two important examples within astrophysics include pulsar timing [35], [50] and large-scale galactic structure studies [34]. In both of these areas, projects have been unable to use Gaussian processes with dense covariance kernels at full resolution of their data because of prohibitive computational scaling. With the arrival of the public *Kepler* mission data – locating tiny exoplanet transit signals superimposed on the stochastic brightness fluctuations of stars – Gaussian processes are an attractive option for the modeling of stellar variability. However, again, the data sets are too large to permit probability calculations with dense kernels at the full resolution of the data with traditional computational techniques. The methods of this paper bring these types of calculations within reach.

The paper is organized as follows. Section 2 reviews some basic facts about Gaussian processes and the resulting formulas encountered in the case of a one-dimensional regression problem. Prediction, marginalization, adaptation of hyperparameters, and existing approximate accelerated methods are also discussed. Section 3 discusses the newly developed matrix factorization for Hierarchical Off-Diagonal Low-Rank (HOLDR) matrices, for which factorization requires only $\mathcal{O}(n \log^2 n)$ work. Subsequent applications of the operator and its inverse scale as $\mathcal{O}(n \log n)$. Many popular covariance functions used for Gaussian processes yield covariance matrices satisfying the HOLDR requirements. While other hierarchical methods could be used for this step, we focus on the HOLDR decomposition because of its simplicity and applicability to a wide range of covariance functions. We would like to emphasize that the algorithm will work for any covariance kernel, but the scaling of the algorithm might not be optimal; for instance, if the covariance kernel has a singularity or is highly oscillatory without damping. Further, in Section 4, we show that the determinant of an HODLR decomposition can be computed in $\mathcal{O}(n \log n)$ operations. Section 5 contains numerical results for our method applied to some standard covariance functions for data embedded in one, two, and three dimensions and to a real data set from the NASA *Kepler* mission. Finally, in Section 6, we summarize our results and discuss other applications of the method, as well as future avenues of research.

The appendix contains a cursory description of the corresponding software packages in C++ and Python which implement the numerical schemes of this work. These open-source software packages have been made available since the time of publication.

2 Gaussian processes and regression

In the past two decades, Gaussian processes have gained popularity in the fields of machine learning and data analysis for their flexibility and robustness. Often cited as a competitive alternative to neural networks because of their rich mathematical and statistical underpinnings, practical use in large-scale problems remains out of reach due to computational complexity. Existing direct computational methods for manipulations involving large-scale covariance matrices require $\mathcal{O}(n^3)$ calculations. This causes regression/prediction, parameter marginalization, and optimization of hyperparameters to be intractable problems. This scaling can be reduced in special cases via several approximation methods, discussed in Section 2.4, however for dense, highly coupled covariance matrices no suitable methods have been proposed.

The following sections contain an overview of regression via Gaussian processes and the large computational tasks that are required at each step. The one-dimensional regression case is discussed for simplicity, but similar formulae for higher dimensions and classification problems are straightforward to derive. In higher dimensions, the corresponding computational methods scale with the same asymptotic complexity, albeit with larger constants. For a thorough treatment of regression using Gaussian processes, see [37], [43].

The canonical linear regression problem we will analyze assumes a model of the form

$$y = f(x) + \epsilon, \tag{2}$$

where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$ is some form of measurement noise. We will enforce the prior distribution of the unknown function f to be a Gaussian process,

$$f \sim \mathcal{GP}(m, k) \tag{3}$$

where $k = k(x, x')$ is some positive definite covariance function, possibly depending on some unknown hyperparameters $\boldsymbol{\theta}$, and $m = m(x)$ is the expected mean of f. The task of fitting hyperparameters is discussed in Sections 2.2 and 2.3 of the paper. Table 1 lists some of

TABLE 1
Covariance kernel functions for Gaussian random process

Name	Covariance function
Ornstein-Uhlenbeck	$\exp\left(-\lvert x-y\rvert\right)$
Gaussian	$\exp\left(-\lvert x-y\rvert^2\right)$
Matérn family	$\sigma^2\frac{\left(\sqrt{2\nu}\lvert x-y\rvert\right)^\nu}{\Gamma(\nu)2^{\nu-1}}K_\nu\left(\sqrt{2\nu}\lvert x-y\rvert\right)$
Rational Quadratic	$\frac{1}{(1+\lvert x-y\rvert^2)^\alpha}$

the frequently used covariance functions for Gaussian processes.

2.1 Prediction

One of the main uses for the previous model (especially in machine learning) is to infer f and predict, with some confidence, $f(\tilde{x})$ for some new input data point $\tilde{x}$. We therefore wish to calculate the conditional distribution $\tilde{y}|\boldsymbol{x},\boldsymbol{y},\tilde{x}$. We will address the nonparametric regression problem, where we assume no analytical form of f, and enforce structure only through the choice of the covariance function k. For the time being, we assume that k is fixed (i.e., hyperparameters are either fixed or absent). Given the data $\boldsymbol{x} = (x_1\ x_2\ \ldots\ x_n)^t$ and $\boldsymbol{y} = (y_1\ y_2\ \ldots\ y_n)^t$, it is easy to show that the conditional distribution of $\boldsymbol{y}$ is given by

$$\boldsymbol{y}|\boldsymbol{x}, f \sim \mathcal{N}\left(\boldsymbol{m}(\boldsymbol{x}),\, \sigma_\epsilon^2 I + K(\boldsymbol{x})\right), \tag{4}$$

where the mean vector and covariance matrix are:

$$\begin{aligned}\boldsymbol{m}(\boldsymbol{x}) &= \left(m(x_1)\ m(x_2)\ \ldots\ m(x_n)\right)^t,\\ K_{ij}(\boldsymbol{x}) &= k(x_i,x_j).\end{aligned} \tag{5}$$

The conditional distribution of $\tilde{y}$ can then be calculated as

$$\tilde{y}|\boldsymbol{x},\boldsymbol{y},\tilde{x} \sim \mathcal{N}(\tilde{f},\tilde{\sigma}^2), \tag{6}$$

with

$$\begin{aligned}\tilde{f} &= \boldsymbol{k}(\tilde{x},\boldsymbol{x})\left(\sigma_\epsilon^2 I + K(\boldsymbol{x})\right)^{-1}\boldsymbol{y},\\ \tilde{\sigma}^2 &= k(\tilde{x},\tilde{x}) - \boldsymbol{k}(\tilde{x},\boldsymbol{x})\left(\sigma_\epsilon^2 I + K(\boldsymbol{x})\right)^{-1}\boldsymbol{k}(\boldsymbol{x},\tilde{x}).\end{aligned} \tag{7}$$

The vector $\boldsymbol{k}(\tilde{x},\boldsymbol{x})$ is the column vector of covariances between $\tilde{x}$ and all the known data points $\boldsymbol{x}$, and $\boldsymbol{k}(\boldsymbol{x},\tilde{x}) = \boldsymbol{k}(\tilde{x},\boldsymbol{x})^t$. See Rasmussen [43] for a derivation of the above formulas. We have therefore reduced the problem of prediction and confidence estimation (in the expected value sense) down to matrix-vector multiplications. For large n, the cost of inverting the matrix $\sigma^2 I + K$ is expensive, with direct methods for dense systems scaling as $\mathcal{O}(n^3)$. A direct algorithm for the rapid inversion of this matrix is one of the main contributions of this paper.

2.2 Hyperparameters and marginalization

As mentioned earlier, often the covariance function k used to model the data $\boldsymbol{x}$, $\boldsymbol{y}$ depends on some set of parameters, $\boldsymbol{\theta}$. For example, in the case of a Gaussian covariance function

$$k(x,x';\boldsymbol{\theta}) = \beta + \frac{1}{\sqrt{2\pi}\sigma}e^{\frac{(x-x')^2}{2\sigma^2}}, \tag{8}$$

the column vector of hyperparameters is given by $\boldsymbol{\theta} = (\beta,\sigma)^t$.

Often hyperparameters correspond to some physically meaningful quantity of the data, for example, a decay rate or some spatial scale. In this case, these parameters are fixed once and for all according to the specific physics or dynamics of the model. On the other hand, hyperparameters may be included for robustness or uncertainty quantification and must be marginalized (integrated) away before the final posterior distribution is calculated. In this case, if prior distributions for these *nuisance* parameters is given by the joint measure $\boldsymbol{\eta}$, then in order to compute the evidence one must compute marginalization integrals of the form

$$p(\boldsymbol{y},\boldsymbol{x}|\boldsymbol{\theta}) \propto \int \frac{1}{(2\pi)^{n/2}|\det(C)|^{1/2}}e^{-\frac{1}{2}(\boldsymbol{x}^t C^{-1}\boldsymbol{x})}\,d\boldsymbol{\eta}, \tag{9}$$

where C is the covariance matrix of the underlying Gaussian process. If there are r nuisance parameters, this is an r-dimensional integral whose numerical integration requires $\mathcal{O}(q^r)$ quadrature nodes, where q is roughly the number of quadrature nodes needed for one-dimensional marginalization. Unless C^{-1} and $\det(C)$ can be calculated rapidly for varying samples of the nuisance parameters, the direct calculation of this integral is not possible. Monte Carlo methods are applicable here, but convergence is only of order one-half. Rapid algorithms for constructing C^{-1} and $\det(C)$ would allow for the *direct* marginalization of nuisance parameters, thereby *directly* constructing the probability of the data, or the marginal evidence.

2.3 Adaptation of hyperparameters

Alternatively, there exist situations in which the hyperparameters $\boldsymbol{\theta}$ do not arise out of physical considerations, but rather one would like to infer them as *best fit parameters*. This entails minimizing some regression norm with respect to the parameters,

$$\min_{\boldsymbol{\theta}} |\boldsymbol{y} - f(\boldsymbol{x})|$$

or rather maximizing a parameter likelihood function (point estimation using a Bayesian framework):

$$\max_{\boldsymbol{\theta}} p(\theta|\boldsymbol{x},\boldsymbol{y},f).$$

In either case, some manner of non-linear optimization must be performed because of the non-linear dependence of *every entry* of the covariance matrix C on the hyperparameters $\boldsymbol{\theta}$.

Regardless of the type of optimization scheme selected (including Monte Carlo-based schemes), several evaluations of the evidence, likelihood, and/or Gaussian regression must be performed – each of which requires

evaluation of the inverse of the covariance matrix, C^{-1}. Unless this inverse can be re-calculated and applied to the data $\boldsymbol{x}$ rapidly, optimizing over all possible $\boldsymbol{\theta}$'s is not a computationally tractable problem. We skip the discussion of various optimization procedures relevant to the adaptation of hyperparameters in Gaussian processes [43] , but only point out that almost all of them require the re-computation of the inverse covariance matrix C^{-1}.

2.4 Accelerated methods

A variety of methods have been proposed to accelerate either the inversion of C, the computation of its determinant, or both. If K is of low-rank, say p, then it is straightforward to compute C^{-1} and $\det(C)$ using the the Sherman-Morrison-Woodbury formula [32], [46], [51] and the Sylvester determinant theorem [1] in $O(p^2n)$ operations. Many of the existing methods for dense matrices are based, in some sense, on finding a suitable low-rank approximation of K:

$$K \approx Q_{n\times p}\, K^s_{p\times p}\, Q^T_{p\times n}\,,$$

where one can think of Q^T as *compressing* the information from K onto a subset of points $\{x_{i1},\ldots,x_{ip}\}$, K^s as the covariance kernel acting on that subset, and Q as *interpolating* the result to the full set of n points $\{x_1, x_2, \ldots, x_n\}$ [40], [47], [48], [49]. Iterative methods can also be applied. These are particularly effective when there is a fast method to compute the necessary matrix-vector products. For Gaussian covariance matrices, this can be accomplished using the fast Gauss transform [27] and its higher-dimensional variants using kd-trees (see, for example, [45], [53]). Alternatively, when $k(x_i, x_j)$ is a convolution kernel, the Fast Fourier Transform (FFT) can be used to accelerate the matrix vector product [16]. As mentioned in the introduction, analysis-based fast algorithms can also be used for specific kernels [15] or treated using the more general "black-box" or "kernel-independent" fast multipole methods [17], [21], [55].

The evaluation of determinants is a somewhat different matter. Taylor series approximations [41] and Monte Carlo methods have been suggested [9], as well as conjugate gradient-type methods combined with trace estimators [15]. For additional approximation methods, see the text [43]. In general, however, it is difficult to obtain accurate values for the determinant in a robust and reliable manner. Thus, the development of a fast, accurate, and direct method is critical in making large-scale Gaussian process modeling feasible.

3 Hierarchical Matrices

A large class of dense matrices, for example, matrices arising out of boundary integral equations [54], radial basis function interpolation [4], kernel density estimation in machine learning, and covariance matrices in statistics and Bayesian inversion [5], [6], can be efficiently represented as data-sparse hierarchical matrices. After a suitable ordering of columns and rows, these matrices can be recursively sub-divided using a tree structure and certain sub-matrices at each level in the tree can be well-represented by low-rank matrices. We refer the readers to [2], [10], [13], [14], [24], [29], [30], [31] for more details on this approach. Depending on the tree structure and low-rank approximation technique, different hierarchical decompositions exist. For instance, the fast multipole method [26] accelerates the calculation of long-range gravitational forces for n-body problems by hierarchically compressing the associated matrix operator using low-rank considerations.

In this article, we will be working with the class of hierarchical matrices known as Hierarchical Off-Diagonal Low-Rank (HODLR) matrices [4], though the ideas extend for other classes of hierarchical matrices as well. As the name suggests, this class of matrices has off-diagonal blocks that are efficiently represented in a recursive fashion using several low-rank matrices. A graphical representation of this class of matrices is shown in Figure 1.

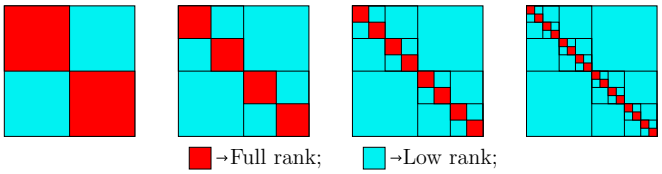

Fig. 1. The HODLR matrix at different levels.

Algebraically, a matrix $K \in \mathbb{R}^{n\times n}$ is termed a two-level HODLR matrix, if it can be written as:

$$K = \begin{bmatrix} K_1^{(1)} & U_1^{(1)}\tilde{K}_{1,2}^{(1)}V_2^{(1)^T} \\ U_2^{(1)}\tilde{K}_{2,1}^{(1)}V_1^{(1)^T} & K_2^{(1)} \end{bmatrix}, \tag{10}$$

with the diagonal blocks given as

$$\begin{aligned} K_1^{(1)} &= \begin{bmatrix} K_1^{(2)} & U_1^{(2)}\tilde{K}_{1,2}^{(2)}V_2^{(2)^T} \\ U_2^{(2)}\tilde{K}_{2,1}^{(2)}V_1^{(2)^T} & K_2^{(2)} \end{bmatrix}, \\ K_2^{(1)} &= \begin{bmatrix} K_3^{(2)} & U_3^{(2)}\tilde{K}_{3,4}^{(2)}V_4^{(2)^T} \\ U_4^{(2)}\tilde{K}_{4,3}^{(2)}V_3^{(2)^T} & K_4^{(2)} \end{bmatrix}, \end{aligned} \tag{11}$$

where $U_i^{(1)}, U_i^{(2)}, V_i^{(1)}, V_i^{(2)}$ are $n \times r$ matrices, $\tilde{K}_{i,j}^{(1)}, \tilde{K}_{i,j}^{(2)}$ are $r \times r$ matrices, and $r \ll n$. In general, all off diagonal blocks of all factors on all levels can be well-represented by a low-rank matrix, i.e., on each level, $U_i^{(k)}, V_i^{(k)}$ are tall and thin matrices and $\tilde{K}_{i,j}^{(k)}$ are small square matrices.

In particular, in what follows we will describe fast direct algorithms for matrix inversion and determinant computation for HODLR matrices in the context of Gaussian processes.

3.1 Fast low-rank approximation of off-diagonal blocks

The first key step is to have a computationally efficient way of obtaining the low-rank factorization of

the off-diagonal blocks. Given any matrix $A \in \mathbb{R}^{m\times n}$, the *optimal* low-rank approximation (in the least-squares sense) is obtained using the singular value decomposition (SVD) [22]. The downside of using the SVD is that the computational cost of direct factorizations scales as $\mathcal{O}(mn\min(m,n))$. For our algorithm to be computationally tractable, we need a fast low-rank factorization. More precisely, we need algorithms that scale as $\mathcal{O}(r^2 n)$ to obtain a rank r factorization of a $n \times n$ matrix. Thankfully, there has recently been tremendous progress in obtaining fast low-rank factorizations of matrices. These techniques can be broadly classified as either *analytic* or *algebraic* techniques.

If the matrix entries are obtained from an analytic function, as is the case for most of the covariance matrices in data analysis, we can rely on approximation theory based analytic techniques like analytic interpolation, multipole expansion, eigenfunction expansion, Taylor series expansions, etc. to obtain a low-rank decomposition. On the other hand, if there is no *a-priori* information of the matrix, then algebraic methods provide an attractive way of computing fast low-rank decompositions. These include techniques like pseudo-skeletal approximations [23], interpolatory decomposition [17], randomized algorithms [19], [36], [52], rank-revealing LU [39], [42], adaptive cross approximation [44] (ACA) (which is a minor variant of partial pivoted LU), and rank-revealing QR [28]. Though analytic techniques can be faster since many operations can be pre-computed, algebraic techniques are attractive for constructing black-box low-rank factorizations. The next section presents the fast matrix factorization of the entire covariance matrix once the low-rank decomposition of the off-diagonal blocks has been obtained using one of the above mentioned techniques. We omit a full discussion of how to actually factorize an HODLR matrix since it would be repeating much of the material in [4].

3.2 HODLR matrix factorization

The overall idea behind the $\mathcal{O}(n\log^2 n)$ factorization of an $n \times n$, κ-level (where $\kappa \sim \log n$) HODLR matrix described in [4] is to factor it as a product of $\kappa+1$ block diagonal matrices,

$$K = K_\kappa\, K_{\kappa-1}\, K_{\kappa-2} \cdots K_1\, K_0, \tag{12}$$

where $K_k \in \mathbb{R}^{n\times n}$ is a block diagonal matrix with 2^k diagonal blocks, each of size $2^{-k}n \times 2^{-k}n$. More importantly, each of these diagonal blocks is a low-rank update to the identity matrix. This factorization relies on the Sherman-Morrison-Woodbury formula [32], [46], [51]. For example, a two-level HODLR matrix described in equations (10) and (11) can be factorized as:

$$\begin{bmatrix} K_1^{(2)} & 0 & 0 & 0 \\ 0 & K_2^{(2)} & 0 & 0 \\ 0 & 0 & K_3^{(2)} & 0 \\ 0 & 0 & 0 & K_4^{(2)} \end{bmatrix} \begin{bmatrix} I_{n/4} & \tilde{K}_{12}^{(2)} & 0 & 0 \\ \tilde{K}_{21}^{(2)} & I_{n/4} & 0 & 0 \\ 0 & 0 & I_{n/4} & \tilde{K}_{34}^{(2)} \\ 0 & 0 & \tilde{K}_{43}^{(2)} & I_{n/4} \end{bmatrix} \begin{bmatrix} I_{n/2} & \tilde{K}_{12}^{(1)} \\ \tilde{K}_{21}^{(1)} & I_{n/2} \end{bmatrix}, \tag{13}$$

where I_m is the $m \times m$ identity matrix, and the matrices $\tilde{K}_{ij}^{(k)}$ are low-rank. Similarly, Figure 2 graphically depicts the factorization of a level 3 HODLR matrix.

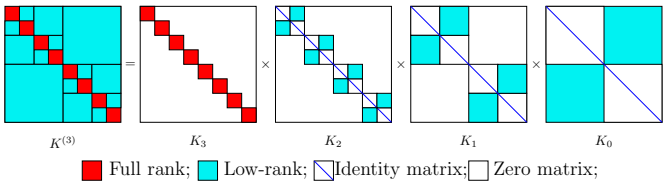

Fig. 2. Factorization of a three level HODLR matrix.

It should be noted that a matrix of the form:

$$\begin{bmatrix} I_n & U_1V_2^T \\ U_2V_1^T & I_n \end{bmatrix} = I_{2n} + \begin{bmatrix} U_1 & 0 \\ 0 & U_2 \end{bmatrix} \begin{bmatrix} 0 & V_2^T \\ V_1^T & 0 \end{bmatrix} \tag{14}$$

is just a low-rank update to the identity matrix. We refer the readers to [4] for more details on the factorization.

4 Determinant Computation

As discussed earlier, once the HODLR factorization has been obtained, Sylvester's determinant theorem [1] enables the computation of the determinant at a cost of $\mathcal{O}(n\log n)$ operations. This new, computationally inexpensive method for direct determinant evaluation enables the efficient *direct* evaluation of probabilities. We now briefly review the algorithm used for determinant evaluation.

Theorem 4.1 (Sylvester's Determinant Theorem)**.** *If $A \in \mathbb{R}^{m\times n}$ and $B \in \mathbb{R}^{n\times m}$, then*

$$\det(I_m + AB) = \det(I_n + BA),$$

where $I_k \in \mathbb{R}^{k\times k}$ is the identity matrix. In particular, for a rank p update to the identity matrix,

$$\det(I_n + U_{n\times p}V_{p\times n}) = \det(I_p + V_{p\times n}U_{n\times p}).$$

Remark 4.2. The computational cost associated with computing the determinant of a rank p update to the identity is $\mathcal{O}(p^2 n)$. The dominant cost is computing the matrix-matrix product $V_{p\times n}U_{n\times p}$.

Furthermore, we recall two basic facts regarding the determinant. First, the determinant of a block diagonal matrix is the product of the determinants of the individual blocks of the matrix. Second, the determinant of a square matrix is completely multiplicative over the set of square matrices, that is to say,

$$\det(A_1A_2\cdots A_n) = \det(A_1)\det(A_2)\cdots\det(A_n). \tag{15}$$

Using the HODLR factorization in equation (12) and these two facts, we have:

$$\det(K) = \det(K_\kappa)\det(K_{\kappa-1})\det(K_{\kappa-2})\cdots \det(K_2)\det(K_1)\det(K_0). \tag{16}$$

Each of the determinants on the right hand side of equation (16) can be computed as a product of the

determinants of diagonal blocks. Each of these diagonal blocks is a low-rank perturbation to the identity, and hence, the determinant can be computed using Sylvester's Determinant Theorem 4.1. It is easy to check that the computational cost for each of the determinants $\det(K_i)$ is $\mathcal{O}(n)$, therefore the total computational cost for obtaining $\det(K)$ is $\mathcal{O}(\kappa n)$, and we recall that $\kappa \sim \log n$.

5 NUMERICAL RESULTS

In this section, we discuss the performance of the previously described algorithm for the inversion and application of covariance matrices C, as well as the calculation of the normalization factor, i.e., the determinant $\det(C)$. Detailed results for n-dimensional datasets $\boldsymbol{x}$, where each x_i is a point in one, two, and three dimensions, are provided for the Gaussian and multiquadric covariance kernels. We also provide benchmarks in one dimension for covariance matrices constructed from exponential, inverse multiquadric, and biharmonic covariance functions. All the proceeding numerical experiments have been run on a MacBook Air with a 1.3GHz Intel Core i5 processor and 4 GB 1600 MHz DDR3 RAM. In all these cases, the matrix entry C_{ij} is given as

$$C_{ij} = \sigma_i^2 \delta_{ij} + k(r_i, r_j)$$

where $k(r_i, r_j)$ is a particular covariance function evaluated at two points, r_i and r_j. In most cases, unless otherwise noted, we have set $\sigma_i = 1$. It should be noted that in certain cases, the matrix K with entries $K_{ij}(\boldsymbol{r}) = k(r_i, r_k)$ might itself be a rank deficient matrix and that this has been exploited in the past to construct fast schemes. However, this is not always the case, for instance, if the covariance function is an exponential. In this situation, the rank of K is in fact full-rank. Even if the matrix K were to be formally rank deficient, in practice, the rank might be very large whereas the ranks of the off-diagonal blocks in the hierarchical structure are very small. Another major advantage of this hierarchical approach is that it is applicable to a wide range of covariance functions and can be used in a black-box, plug-and-play, fashion.

5.1 Gaussian covariance

Here the covariance function, and the corresponding entry of K, is given as

$$k(r_i, r_j) = \exp\left(-|r_i - r_j|^2\right), \tag{17}$$

where r_i, r_j are points in one, two, or three dimensions as noted in the individual tables. In each of the following tables, timings are provided for the assembly, factorization, and inversion of the covariance matrix (denoted by the columns *Assembly*, *Factor*, and *Solve*). The time for computing the determinant of the $n \times n$ matrix is given in the *Det.* column, and the error provided is approximately the relative l_2 precision in the solution to a test problem $C\boldsymbol{x} = \boldsymbol{b}$, where $\boldsymbol{b}$ was generated *a priori* from the known vector $\boldsymbol{x}$. These results are reported in Tables 2, 3, and 4. Scaling of the algorithm for data embedded in one, two, and three dimensions is compared with the direct calculation in Figure 3.

TABLE 2
Timings for one-dimensional Gaussian covariance functions. The matrix entry is given as $C_{ij} = \delta_{ij} + \exp\left(-|r_i - r_j|^2\right)$, where r_i are random uniformly distributed points in the interval $[-3, 3]$.

	Time taken in seconds				
n	Assembly	Factor	Solve	Det.	Error
10^4	0.12	0.11	0.008	0.01	10^{-13}
2×10^4	0.15	0.23	0.016	0.03	10^{-13}
5×10^4	0.47	0.71	0.036	0.12	10^{-12}
10^5	1.24	1.46	0.052	0.24	10^{-12}
2×10^5	2.14	3.12	0.121	0.39	10^{-13}
5×10^5	6.13	10.2	0.388	0.56	10^{-12}
10^6	14.1	23.2	0.834	1.52	10^{-12}

TABLE 3
Timings for two-dimensional Gaussian covariance functions. The matrix entry is given as $C_{ij} = \delta_{ij} + \exp\left(-|r_i - r_j|^2\right)$, where r_i are random uniformly distributed points in the square $[-3, 3] \times [-3, 3]$.

	Time taken in seconds				
n	Assembly	Factor	Solve	Det.	Error
10^4	0.56	0.50	0.018	0.03	10^{-13}
2×10^4	1.16	0.99	0.028	0.05	10^{-13}
5×10^4	2.74	2.44	0.067	0.12	10^{-13}
10^5	5.43	5.08	0.165	0.23	10^{-12}
2×10^5	12.4	14.4	0.485	0.44	10^{-12}
5×10^5	31.7	37.3	1.33	1.17	10^{-12}
10^6	70.8	79.2	3.15	2.24	10^{-12}

TABLE 4
Timings for three-dimensional Gaussian covariance functions. The matrix entry is given as $C_{ij} = \delta_{ij} + \exp\left(-|r_i - r_j|^2\right)$, where r_i are random uniformly distributed points in the cube $[-3, 3] \times [-3, 3] \times [-3, 3]$.

	Time taken in seconds				
n	Assembly	Factor	Solve	Det.	Error
10^4	15.4	17.3	0.113	0.91	10^{-12}
2×10^4	30.9	33.1	0.224	1.06	10^{-12}
5×10^4	75.5	76.3	0.434	1.68	10^{-11}
10^6	149	166	0.923	3.11	10^{-11}

5.2 Multiquadric covariance matrices

Covariance functions of the form

$$k(r_i, r_j) = \sqrt{(1 + |r_i - r_j|^2)} \tag{18}$$

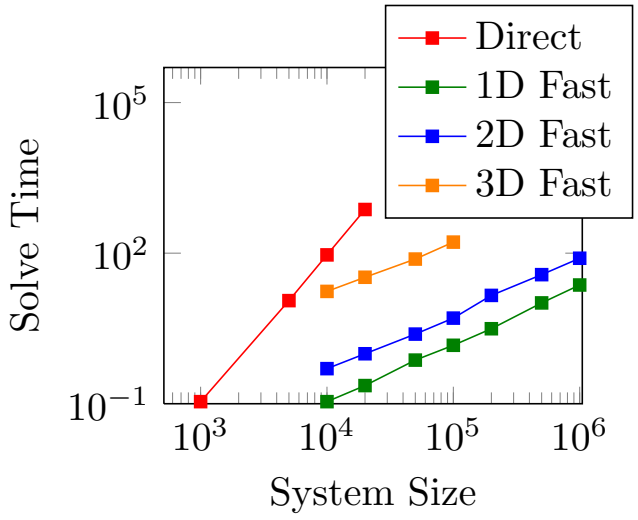

Fig. 3. Comparison of time required to factorize the covariance matrix in the case of a Gaussian covariance kernel in one, two, and three dimensions. The conventional direct calculation is independent of dimension but scales as $\mathcal{O}(n^3)$.

are known as multiquadric covariance functions, one class of frequently used radial basis functions. Analogous numerical results are presented below in one, two, and three dimensions as were in the previous section for the Gaussian covariance function. Tables 5, 6, and 7 contain timings and Figure 4 shows scaling results.

TABLE 5
Timings for one-dimensional multiquadric covariance functions. The matrix entry is given as $C_{ij} = \delta_{ij} + \sqrt{1 + |r_i - r_j|}$, where r_i are random uniformly distributed points in the interval $[-3, 3]$.

	Time taken in seconds				
n	Assembly	Factor	Solve	Det.	Error
10^4	0.08	0.13	0.006	0.02	10^{-13}
2×10^4	0.11	0.22	0.011	0.03	10^{-13}
5×10^4	0.34	0.65	0.030	0.13	10^{-13}
10^5	0.85	1.44	0.059	0.22	10^{-12}
2×10^5	1.56	3.12	0.147	0.44	10^{-13}
5×10^5	4.72	8.33	0.363	0.94	10^{-12}
10^6	10.9	17.1	0.814	1.94	10^{-12}

5.3 Exponential covariance

Covariances functions of the form

$$k(r_i, r_j) = \exp(-|r_i - r_j|) \tag{19}$$

are known as exponential covariance functions. Analogous numerical results are presented in Table 8 for one dimension as were in the previous sections for Gaussian and multiquadric covariance functions. Figure 5 compares scaling for various kernels.

TABLE 6
Timings for two-dimensional multiquadric covariance functions. The matrix entry is given as $C_{ij} = \delta_{ij} + \sqrt{1 + |r_i - r_j|^2}$, where r_i are random uniformly distributed points in the square $[-3, 3] \times [-3, 3]$.

	Time taken in seconds				
n	Assembly	Factor	Solve	Det.	Error
10^4	0.77	0.86	0.022	0.04	10^{-13}
2×10^4	1.41	1.42	0.042	0.06	10^{-13}
5×10^4	3.31	3.43	0.082	0.15	10^{-12}
10^5	6.54	6.95	0.177	0.31	10^{-12}
2×10^5	14.1	15.9	0.395	0.59	10^{-11}
5×10^5	38.2	42.1	1.12	1.69	10^{-11}
10^6	79.9	90.3	2.38	3.39	10^{-11}

TABLE 7
Timings for three-dimensional multiquadric covariance functions. The matrix entry is given as $C_{ij} = \delta_{ij} + \sqrt{1 + |r_i - r_j|^2}$, where r_i are random uniformly distributed points in the cube $[-3, 3] \times [-3, 3] \times [-3, 3]$.

	Time taken in seconds				
n	Assembly	Factor	Solve	Det.	Error
10^4	19.0	23.2	0.135	1.32	10^{-12}
2×10^4	38.7	45.1	0.276	1.65	10^{-11}
5×10^4	87.8	97.8	0.578	2.28	10^{-11}
10^5	164	195	1.24	3.84	10^{-10}

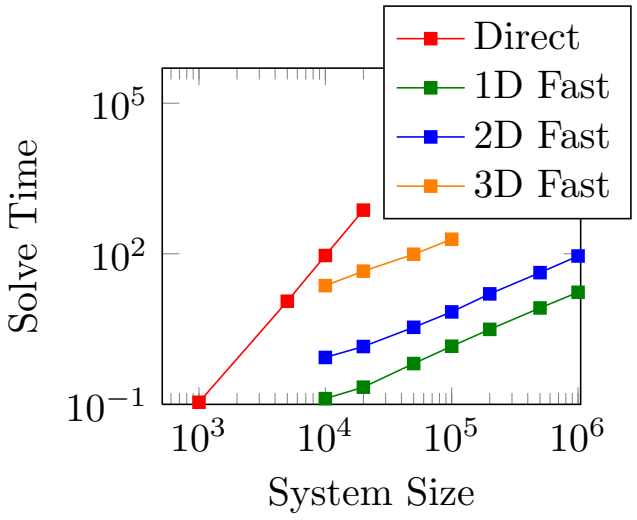

Fig. 4. Comparison of time required to factorize the covariance matrix in the case of a multiquadric covariance kernel in one, two, and three dimensions.

TABLE 8
Timings for one-dimensional exponential covariance functions. The matrix entry is given as $C_{ij} = \delta_{ij} + \exp(-|r_i - r_j|)$, where r_i are random uniformly distributed points in the interval $[-3, 3]$.

	Time taken in seconds				
n	Assembly	Factor	Solve	Det.	Error
10^4	0.13	0.06	0.003	0.02	10^{-13}
2×10^4	0.23	0.11	0.008	0.03	10^{-13}
5×10^4	0.64	0.32	0.020	0.10	10^{-12}
10^5	1.41	0.70	0.039	0.23	10^{-13}
2×10^5	2.86	1.42	0.076	0.42	10^{-12}
5×10^5	8.63	3.47	0.258	0.67	10^{-12}
10^6	18.8	8.05	0.636	1.35	10^{-12}

5.4 Inverse Multiquadric and Biharmonic

The inverse multiquadric and biharmonic kernel (also known as the thin plane spline) are frequently used in radial basis function interpolation and *kriging* in geostatistics. These kernels are given by the formulae

$$k_1(r_i, r_j) = \frac{1}{\sqrt{1 + |r_i - r_j|^2}}, \qquad k_2(r_i, r_j) = |r_i - r_j|^2 \log |r_i - r_j|, \tag{20}$$

respectively. Timing results are presented in Tables 9 and 10, and comparison with the exponential kernel is shown in Figure 5.

TABLE 9
Timings for one-dimensional inverse multiquadric covariance functions. The matrix entry is given as $C_{ij} = \delta_{ij} + 1/\sqrt{1 + |r_i - r_j|^2}$, where r_i are random uniformly distributed points in the interval $[-3, 3]$.

	Time taken in seconds				
n	Assembly	Factor	Solve	Det.	Error
10^4	0.11	0.13	0.006	0.02	10^{-13}
2×10^4	0.17	0.29	0.017	0.04	10^{-13}
5×10^4	0.47	0.84	0.037	0.12	10^{-12}
10^5	1.07	1.58	0.072	0.21	10^{-12}
2×10^5	2.18	3.49	0.158	0.44	10^{-12}
5×10^5	6.43	11.8	0.496	0.71	10^{-11}
10^6	14.2	26.8	1.02	1.49	10^{-11}

5.5 A real data set: *Kepler* Spacecraft light curve time series

The *Kepler* Mission is a NASA mission designed to find transiting exoplanets – planets orbiting stars beyond the solar system – by measuring the brightness of approximately 150,000 stars as a function of time. The main *Kepler* Mission collected time-series data (stellar brightness as a function of time) for all these stars with half-hour exposures for roughly four years. At the time of publication, approximately 3850 exoplanet candidates have been discovered using this dataset (a real-time update of this statistic can be found at exoplanetarchive.ipac.caltech.edu).

TABLE 10
Timings for one-dimensional biharmonic covariance function. The matrix entry is given as $C_{ij} = 2\delta_{ij} + |r_i - r_j|^2 \log(|r_i - r_j|)$, where r_i are random uniformly distributed points in the interval $[-3, 3]$.

	Time taken in seconds				
n	Assembly	Factor	Solve	Det.	Error
10^4	0.28	0.31	0.015	0.03	10^{-12}
2×10^4	0.61	0.59	0.028	0.06	10^{-12}
5×10^4	1.62	1.68	0.067	0.15	10^{-12}
10^5	3.61	3.93	0.123	0.34	10^{-12}
2×10^5	8.03	10.7	0.236	0.65	10^{-12}
5×10^5	26.7	31.2	0.632	1.41	10^{-11}
10^6	51.3	81.9	1.28	3.40	10^{-11}

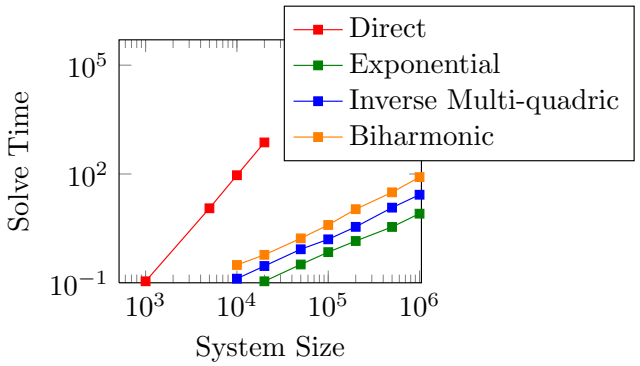

Fig. 5. Comparison of time required to factorize a variety of covariance matrices arising from exponential, inverse multiquadric, and biharmonic covariance kernels for one-dimensional data.

We now apply the previously described hierarchical matrix decomposition machinery to the problem of regressing light-intensity data upon timestamps obtained from the *Kepler* Mission. The regression is performed as in Section 2, that is, assuming a model of a Gaussian process. This is a useful model because the data is "contaminated" by stochastic stellar variability and other noisy processes without a physical model. Gibson et al. [20] demonstrated the use of Gaussian processes as a model for the stochastic noise sources in a photometric light curve, but the standard tools used for evaluating the model do not scale to datasets as large as a *Kepler* Spacecraft light curve.

There is one main reason why our HODLR-based technique is necessary for this study. The Gaussian process is only an effective model *of the noise*, so we would like to *marginalize* over the hyperparameters describing the covariance function. This means that we must both compute both the determinant and apply the inverse of the covariance matrix millions of times inside a Monte Carlo sampling routine. As we demonstrate below, the HODLR method provides a computationally tractable solution to

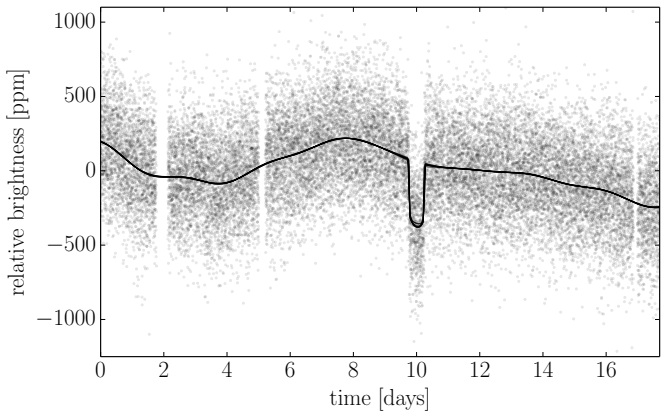

Fig. 6. *Kepler* data on KIC 10593626 with a synthetic injected transit at 10 days and posterior samples from an MCMC analysis with nine parameters (7 physical and 2 hyperparameters for the noise model). The lines are the predictive mean functions for ten posterior samples from the Markov chain.

both of these issues for large datasets obtained from the *Kepler* Mission.

Table 11 reports the computation costs of evaluating a realistic Gaussian process noise model on a real *Kepler* star. The catalog ID for the star is KIC 12365184 and the dataset is a *short cadence* light curve with samples every minute. Figure 6 shows the results of an MCMC analysis on a synthetic transit injected into the light curve of KIC 10593626. The sampled parameters are seven physical parameters describing the star and the orbit, and two hyperparameters (α and s) specifying the covariance function k and matrix entries:

$$C_{ij}(t_i,\, t_j;\, \alpha,\, s) = \sigma_i^2\, \delta_{ij} + \alpha\, \exp\left(-\frac{(t_i - t_j)^2}{2\, s}\right) \qquad (21)$$

where the σ_i^2 are estimates of the measurement uncertainties. The points in Figure 6 are the data and the lines are the predictive mean functions for ten random posterior samples (see [43] for a description of the predictive mean computation). This analysis would not have been computationally tractable without the methods described in this article.

TABLE 11
The computational costs of evaluating a Gaussian process noise model on *Kepler* short cadence data for star KIC 12365184.

	Time taken in seconds		
n	Assembly & Factor	Solve	Determinant
100	0.001	0.00005	0.0003
500	0.011	0.00023	0.0008
1000	0.028	0.00058	0.0018
5000	0.238	0.00575	0.0123
10000	0.583	0.00915	0.0224
47911	4.35	0.0476	0.126

6 Conclusions

In this paper, we have presented a fast, accurate, and nearly optimal hierarchical direct linear algebraic algorithms for computing determinants, inverses, and matrix-vector products involving covariances matrices encountered when using Gaussian processes. Similar matrices appear in problems of classification and prediction; our method carries over and applies equally well to these problems. Previous attempts at accelerating these calculations (inversion and determinant calculation) relied on either sacrificing fidelity in the covariance kernel (e.g., thresholding) or paying the computational penalty of dealing with dense, full-rank covariance matrices. Our HODLR-based algorithm obviates the need for this compromise.

The source code for the algorithm has been made available on GitHub. The HODLR package for solving linear systems and computing determinants is available at https://github.com/sivaramambikasaran/HODLR [3] and the Python Gaussian process package [18], *george*, has been made available at https://github.com/dfm/george. See the appendix for information regarding these software packages.

In its present form, our method degrades in performance when the n-dimensional data has a covariance function based on points in $\mathbb{R}^d$ with $d > 3$, as well as when the covariance function is oscillatory. Extensions of our approach to these cases is a subject of current research. We are also investigating high-dimensional anisotropic quadratures for marginalization and moment computation.

Appendix

This publication also serves to formally announce the release of an HODLR-based fast Gaussian processes package, *george*, written in Python, available on GitHub, and supported by D. Foreman-Mackey. This package supplies a Python wrapper and interface to the C++ HODLR software [3]. Both packages are open source, the HODLR package is released under the MPL2.0 license and *george* [18] is released under the MIT license. Details on using these packages are available at their respective online repositories.

Acknowledgments

The authors would like to thank Iain Murray for several useful and detailed discussions.

References

[1] A. G. Akritas, E. K. Akritas, and G. I. Malaschonok. Various proofs of Sylvester's (determinant) identity. Mathematics and Computers in Simulation, 42(4):585–593, 1996.

[2] S. Ambikasaran. Fast Algorithms for Dense Numerical Linear Algebra and Applications. PhD thesis, Stanford University, 2013.

[3] S. Ambikasaran. A fast direct solver for dense linear systems. https://github.com/sivaramambikasaran/HODLR, 2013.

[4] S. Ambikasaran and E. Darve. An $\mathcal{O}(n \log n)$ Fast Direct Solver for Partial Hierarchically Semi-Separable Matrices. Journal of Scientific Computing, 57(3):477–501, 2013.

[5] S. Ambikasaran, J. Y. Li, P. K. Kitanidis, and E. Darve. Large-scale stochastic linear inversion using hierarchical matrices. Computational Geosciences, 17(6):913–927, 2013.

[6] S. Ambikasaran, A. K. Saibaba, E. F. Darve, and P. K. Kitanidis. Fast Algorithms for Bayesian Inversion. In Computational Challenges in the Geosciences, pages 101–142. Springer, 2013.

[7] A. Aminfar, S. Ambikasaran, and E. Darve. A fast block low-rank dense solver with applications to finite-element matrices. arXiv preprint arXiv:1403.5337 [cs-NA], 2014.

[8] M. Anitescu, J. Chen, and L. Wang. A matrix-free approach for solving the parametric gaussian process maximum likelihood problem. SIAM J. Sci. Comput., 34:A240–A262, 2012.

[9] R. P. Barry and R. K. Pace. Monte carlo estimates of the log determinant of large sparse matrices. Linear Algebra and its Applications, 289:41–54, 1999.

[10] S. Börm, L. Grasedyck, and W. Hackbusch. Hierarchical matrices. Lecture notes, 21, 2003.

[11] R. P. Brent, J.-A. H. Osborn, and W. D. Smith. Bounds on determinants of perturbed diagonal matrices. arxiv.org/abs/1401.7084, 2013.

[12] K. Chalupka, C. K. I. Williams, and I. Murray. A framework for evaluating approximation methods for gaussian process regression. J. Mach. Learn. Res., 14:333–350, 2013.

[13] S. Chandrasekaran, P. Dewilde, M. Gu, W. Lyons, and T. Pals. A fast solver for hss representations via sparse matrices. SIAM Journal on Matrix Analysis and Applications, 29(1):67–81, 2006.

[14] S. Chandrasekaran, M. Gu, and T. Pals. A fast ULV decomposition solver for hierarchically semiseparable representations. SIAM Journal on Matrix Analysis and Applications, 28(3):603–622, 2006.

[15] J. Chen, L. Wang, and M. Anitescu. A fast summation tree code for matérn kernel. SIAM J. Sci. Comput., 36:A289–A309, 2013.

[16] C. Dietrich and G. Newsam. Fast and exact simulation of stationary gaussian processes through circulant embedding of the covariance matrix. SIAM Journal on Scientific Computing, 18:1088, 1997.

[17] W. Fong and E. Darve. The black-box fast multipole method. Journal of Computational Physics, 228(23):8712–8725, 2009.

[18] D. Foreman Mackey. Fast gaussian processes for regression. https://github.com/dfm/george, 2014.

[19] A. Frieze, R. Kannan, and S. Vempala. Fast Monte-Carlo algorithms for finding low-rank approximations. Journal of the ACM (JACM), 51(6):1025–1041, 2004.

[20] N. P. Gibson, S. Aigrain, S. Roberts, T. M. Evans, M. Osborne, and F. Pont. A gaussian process framework for modelling instrumental systematics: application to transmission spectroscopy. MNRAS, 419:2683–2694, jan 2012.

[21] Z. Gimbutas and V. Rokhlin. A generalized fast multipole method for nonoscillatory kernels. SIAM Journal on Scientific Computing, 24(3):796–817, 2003.

[22] G. Golub and C. Van Loan. Matrix computations, volume 3. Johns Hopkins Univ Press, 1996.

[23] S. Goreinov, E. Tyrtyshnikov, and N. Zamarashkin. A theory of pseudoskeleton approximations. Linear Algebra and its Applications, 261(1-3):1–21, 1997.

[24] L. Grasedyck and W. Hackbusch. Construction and arithmetics of $\mathcal{H}$-matrices. Computing, 70(4):295–334, 2003.

[25] L. Greengard, D. Gueyffier, P.-G. Martinsson, and V. Rokhlin. Fast direct solvers for integral equations in complex three-dimensional domains. Acta Numerica, 18(1):243–275, 2009.

[26] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. Journal of Computational Physics, 73(2):325–348, 1987.

[27] L. Greengard and J. Strain. The fast gauss transform. SIAM J. Sci. Stat. Comput., 12:79–94, 1991.

[28] M. Gu and S. Eisenstat. Efficient algorithms for computing a strong rank-revealing QR factorization. Society, 17(4):848–869, 1996.

[29] W. Hackbusch. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. part i: Introduction to $\mathcal{H}$-matrices. Computing, 62(2):89–108, 1999.

[30] W. Hackbusch and S. Börm. Data-sparse approximation by adaptive $\mathcal{H}^2$-matrices. Computing, 69(1):1–35, 2002.

[31] W. Hackbusch and B. N. Khoromskij. A sparse $\mathcal{H}$-matrix arithmetic. Computing, 64(1):21–47, 2000.

[32] W. Hager. Updating the inverse of a matrix. SIAM review, pages 221–239, 1989.

[33] K. L. Ho and L. Greengard. A fast direct solver for structured linear systems by recursive skeletonization. SIAM Journal on Scientific Computing, 34(5):2507–2532, 2012.

[34] J. Jasche and B. D. Wandelt. Methods for bayesian power spectrum inference with galaxy surveys. Astrophys. J., 779:15, 2013.

[35] L. Lentati, P. Alexander, M. P. Hobson, S. Taylor, J. Gair, S. T. Balan, and R. van Haasteren. Hyper-efficient model-independent bayesian method for the analysis of pulsar timing data. Phys. Rev. D, 87:104021, May 2013.

[36] E. Liberty, F. Woolfe, P.-G. Martinsson, V. Rokhlin, and M. Tygert. Randomized algorithms for the low-rank approximation of matrices. Proceedings of the National Academy of Sciences, 104(51):20167, 2007.

[37] D. J. MacKay. Introduction to gaussian processes. NATO ASI Series F Computer and Systems Sciences, 168:133–166, 1998.

[38] P.-G. Martinsson and V. Rokhlin. A fast direct solver for boundary integral equations in two dimensions. Journal of Computational Physics, 205(1):1–23, 2005.

[39] L. Miranian and M. Gu. Strong rank revealing LU factorizations. Linear Algebra and its Applications, 367:1–16, 2003.

[40] J. Q. nonero Candela and C. E. Rasmussen. A unifying view of sparse approximate gaussian process regression. J. Mach. Learn. Res., 6:1939–1959, 2005.

[41] R. K. Pace and J. P. LeSage. Chebyshev approximation of log-determinants of spatial weight matrices. Comput. Stat. & Data Anal., 45:179–196, 2004.

[42] C.-T. Pan. On the existence and computation of rank-revealing LU factorizations. Linear Algebra and its Applications, 316(1):199–222, 2000.

[43] C. E. Rasmussen and C. Williams. Gaussian processes for machine learning. MIT Press, 2006.

[44] S. Rjasanow. Adaptive cross approximation of dense matrices. IABEM 2002, International Association for Boundary Element Methods, 2002.

[45] Y. Shen, A. Ng, and M. Seeger. Fast gaussian process regression using kd-trees. In Y. Weiss, Schölkopf, and J. Platt, editors, Advances in Neural Information Processing Systems 18, pages 1225–1232. MIT Press, Cambridge, MA, 2006.

[46] J. Sherman and W. J. Morrison. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. The Annals of Mathematical Statistics, pages 124–127, 1950.

[47] A. J. Smola and P. L. Bartlett. Sparse greedy gaussian process regression. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, Advances in Neural Information Processing Systems 13, pages 619–625. MIT Press, Cambridge, MA, 2001.

[48] E. Snelson and Z. Ghahramani. Local and global sparse gaussian process approximations. In M. Meila and X. Shen, editors, Artifical Intelligence and Statistics (AISTATS 11), volume 11, pages 524–531, 2007.

[49] E. L. Snelson. Flexible and efficient Gaussian process models for machine learning. PhD thesis, Gatsby Computational Neuroscience Unit, University College London, 2007.

[50] R. van Haasteren. Accelerating pulsar timing data analysis. Monthly Notices of the Royal Astronomical Society, 429(1):55–62, 2013.

[51] M. A. Woodbury. Inverting modified matrices. Statistical Research Group, Memo. Rep. no. 42, Princeton University, 1950.

[52] F. Woolfe, E. Liberty, V. Rokhlin, and M. Tygert. A fast randomized algorithm for the approximation of matrices. Applied and Computational Harmonic Analysis, 25(3):335–366, 2008.

[53] C. Yang, R. Duraiswami, and L. Davis. Efficient kernel machines using the improved fast gauss transform. In L. K. Saul, Y. Weiss, and L. Bottou, editors, Advances in Neural Information Processing Systems 17, pages 1561–1568. MIT Press, Cambridge, MA, 2005.

[54] L. Ying. Fast algorithms for boundary integral equations. In Multiscale Modeling and Simulation in Science, pages 139–193. Springer, 2009.

[55] L. Ying, G. Biros, and D. Zorin. A kernel-independent adaptive fast multipole algorithm in two and three dimensions. Journal of Computational Physics, 196(2):591–626, 2004.